The ESPSRC's approach to an ecosystem of services for science support within the SKA Regional Centres Network

Manuel Parra-Royón, Lourdes Verdes-Montenegro, Susana Sánchez-Expósito, Julian Garrido, Laura Darriba, Javier Moldón, María Ángeles Mendoza, Jesús Sánchez-Castañeda

Instituto de Astrofísica de Andalucía - CSIC

IBERGRID 2023
Benasque, Huesca, Spain

Table of content

Introduction to SKAO

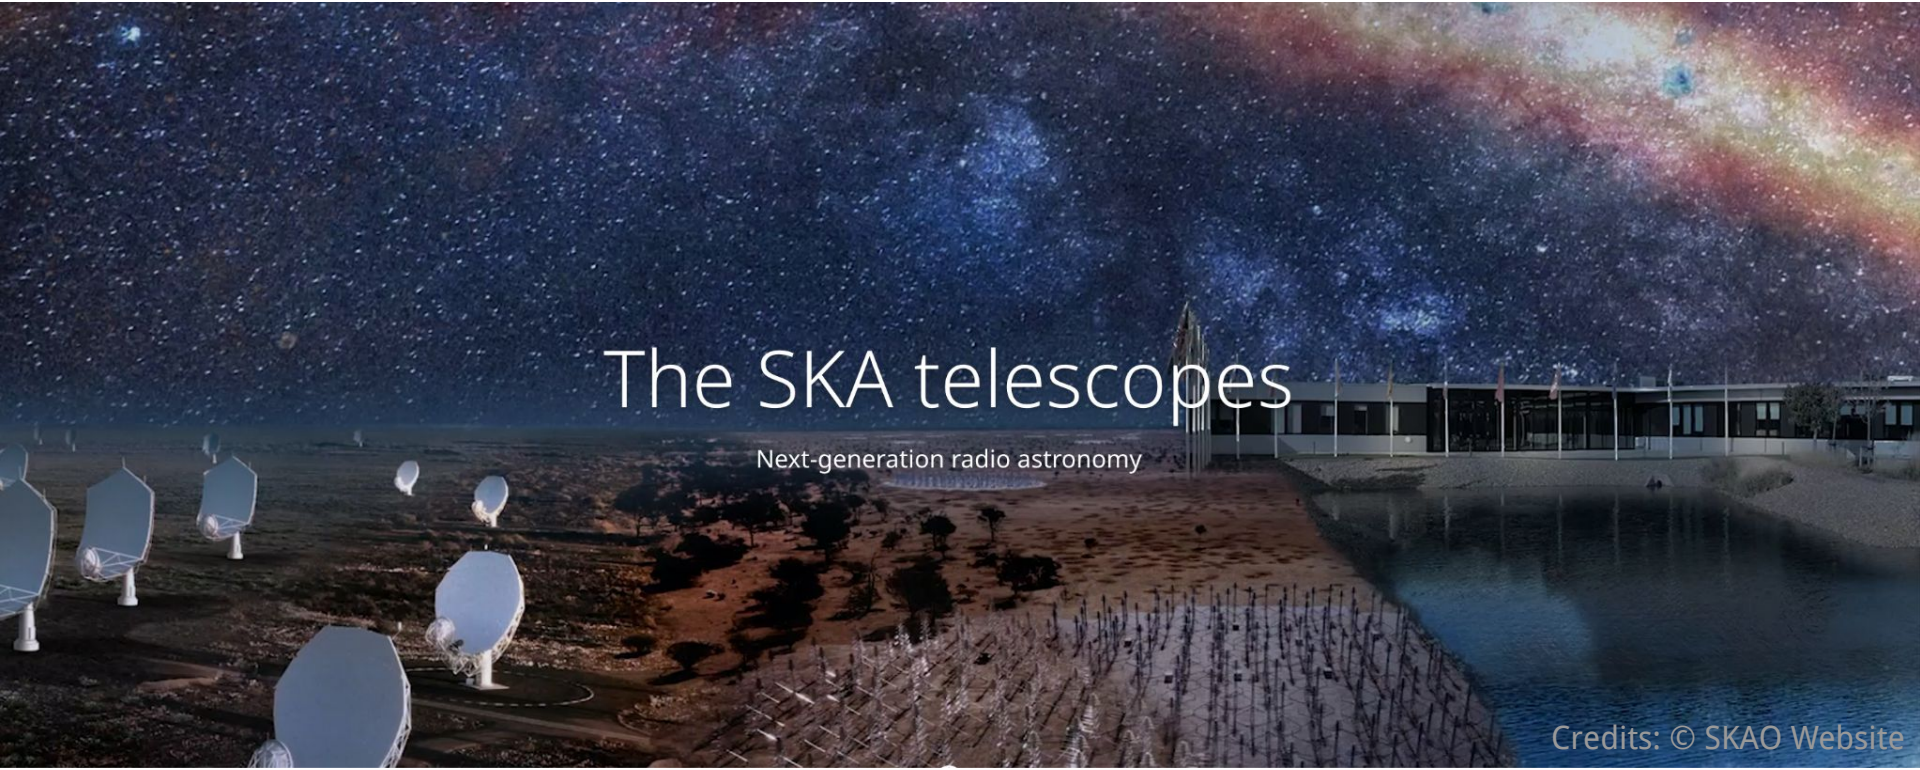

Introduction to SKAO

SKAO is a global organisation responsible for the **development and operation** of the Square Kilometre Array Telescope (**SKA**).

SKAO represents one of the largest and most ambitious scientific projects, aiming to **explore the universe** through radio astronomy.

SKAO and its science will tackle the open key questions in **Astrophysics**, **Astrobiology** and **Fundamental Physics**.

Introduction to SKAO

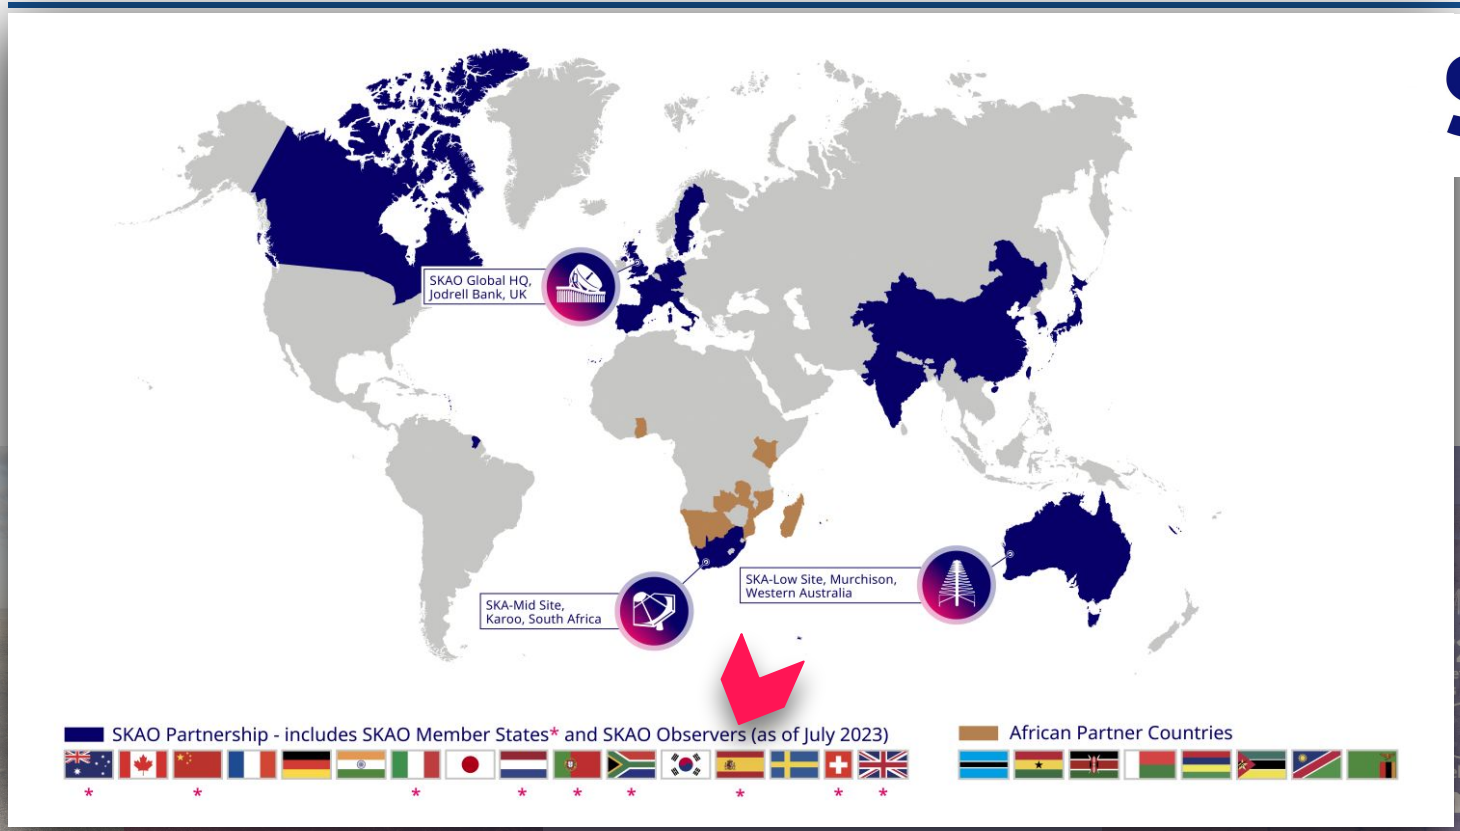

Credits: © SKAO

Introduction to SKAO: The Spanish coordination

- **2011:** Official kick-off with “Feasibility study of the Spanish participation in the SKA”
- **2018/19 + 2020/21 + 2023/2024:** Excellence Network RED-SKA (PI: LVM)

  BSC, CAB, CEFCA, IAC, ICE, IFCA, IFIC, OAN, UGR, UMA, UPM, USAL, USC, UV
- **June 2018:** Spain became a Member of the SKA Organisation
- **2021- :** Coordination of the SKA-Spain participation directly funded by Ministry
- **2021:** Spain’s membership in the SKA Observatory approved.

  Accession process: signed by Council of State, very close final stage - Council of Ministers

Coordinator of the Spanish participation in SKAO
Lourdes Verdes-Montenegro Atalaya (LVM)

SKAO

The SKA Regional Centres Network (SRCNet)

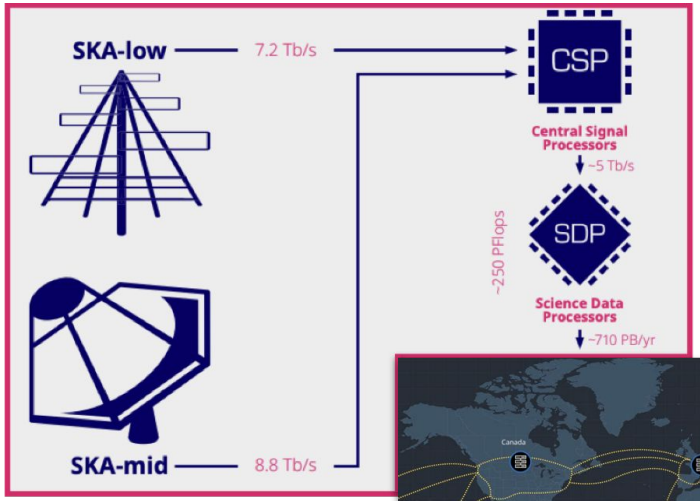

SKA Regional Centres Network

15 on-going national initiatives

SRCNet
Technological challenges

Big Data, e-Science, Data management, Remote visualisation, Science Analysis Algorithms, AI, End-to-End reproducibility, Green computing

The core of the SKA Science

The Spanish SRC: **esp**SRC

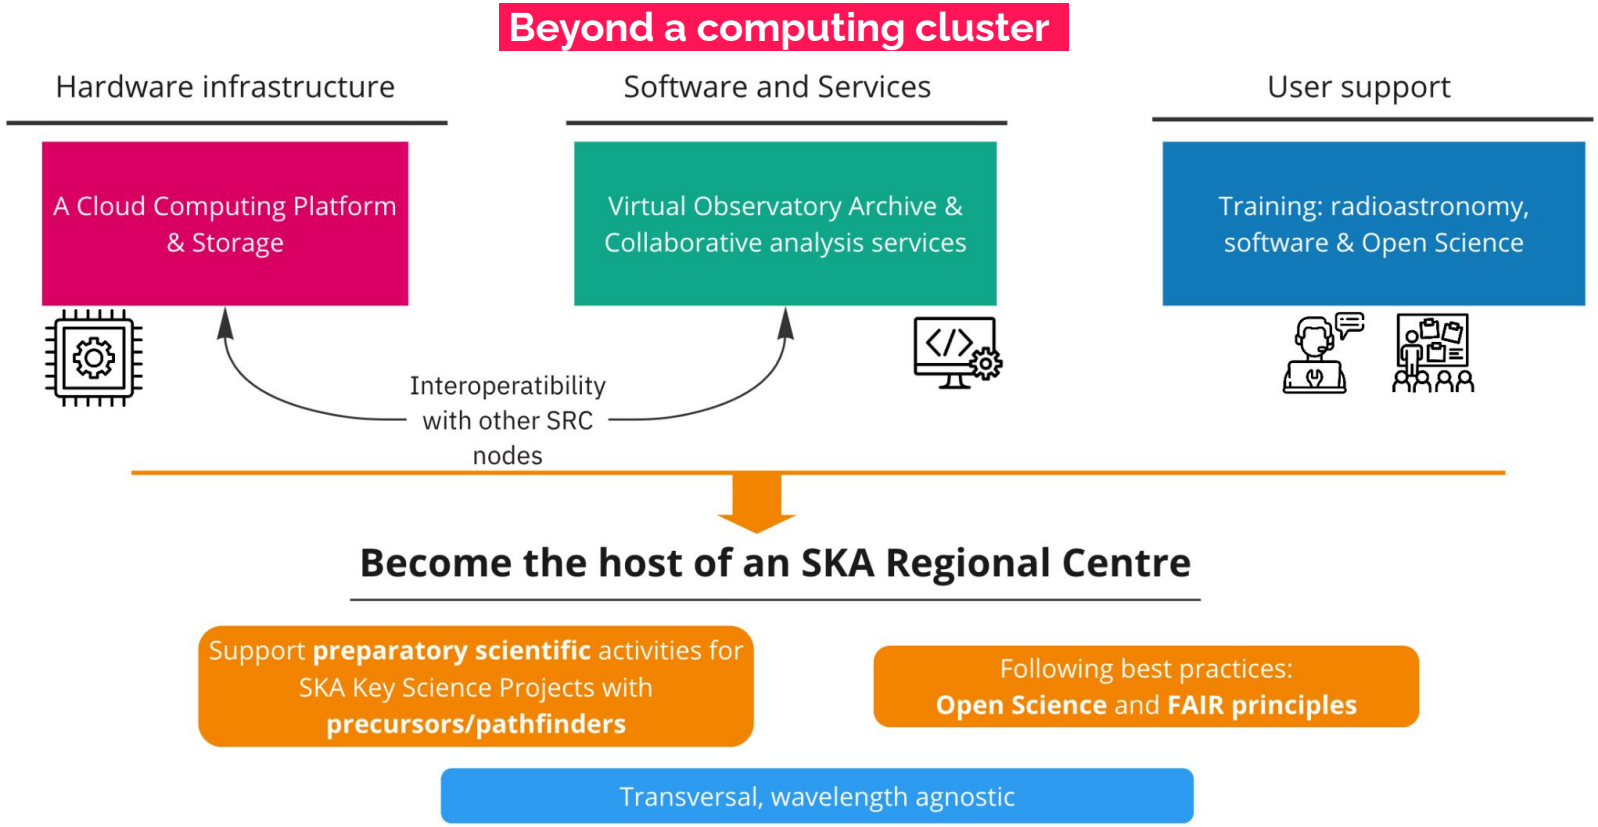

The Spanish SRC: **esp**SRC team

Susana Sánchez
espSRC Technical lead

Laura Darriba
Project manager, data science, user support

Manuel Parra
Cloud computing engineer, Distributed computing

Jesús Sánchez
Sys admin, Open Stack, user support

Javier Moldón
Scientific operations, software deployment, radio astronomer

M. Ángeles Mendoza
Science domain services engineer

Julián Garrido
Science archive, VO standards

Lourdes Verdes-Montenegro
Coordinator of the Spanish participation in the SKAO, espSRC coordinator, radio astronomer

The **esp**SRC computing infrastructure model

Credits: © AMIGA Group

Computing nodes: **Dell PE R640**

200 CPU Cores
4 nodes with 368 GBytes of RAM
1 node with 1 TByte of RAM
2.5 TB of RAM

Storage nodes: **Dell PE R740 XD**

Credits: © Dell

8 nodes,
~ 1.5 PBytes of raw storage

The **esp**SRC computing infrastructure model

Computing nodes

4 Dell PowerEdge R640	CPU: 2 Intel Xeon 6230 (20C/40T) RAM: 24× 16 GB RAM DDR4 2933MT/s moduleswith a total of 384 GB Disks: 2 × 7.68 TB SAS 12 Gbps SSDs Network: Broadcom 5720 with 4 1 GbE ports and Mellanox ConnectX-5 EX with 2 100 GbE QSFP28
1 Dell PowerEdge R640	CPU: 2 Intel Xeon 6230 (20C/40T) RAM: 16 × 64 GB RAM DDR4 2933MT/s modules with a total of 1024 GB Disks: 2 × 7.68 TB SAS 12 Gbps SSDs Network: Broadcom 5720 with 4 1 GbE ports and Mellanox ConnectX-5 EX with 2 100 GbE QSFP28

es of RAM
RAM

Storage

RAW

The **esp**SRC computing infrastructure model

Storage nodes	4 Dell PowerEdge R740xd	CPU: 2 Intel Xeon 6230 (20C/40T) RAM: 16 × 16 GB RAM DDR4 2933 MT/s modules with a total of 256GB Disks: 32 × 7.68 TB SAS 12 Gbps SSDs Network: Broadcom 5720 with 4 1 GbE ports and Mellanox ConnectX-5 EX with 2 100 GbE QSFP28
Storage	4 Dell PowerEdge R740xd	CPU: 2× Intel Xeon 6230 (20C/40T) RAM: 16 × 16 GB RAM DDR4 2933 MT/s modules with a total of 256 GB RAM Disks: 32 × 3.84 TB SAS 12 Gbps SSDs Network: Broadcom 5720 with 4 × 1 GbE ports and Mellanox ConnectX-5 EX with 2 × 100 GbE QSFP28

The **esp**SRC computing infrastructure model

openstack.

Early
2020

Mid
2022

Fall
2023

The **esp**SRC computing infrastructure model

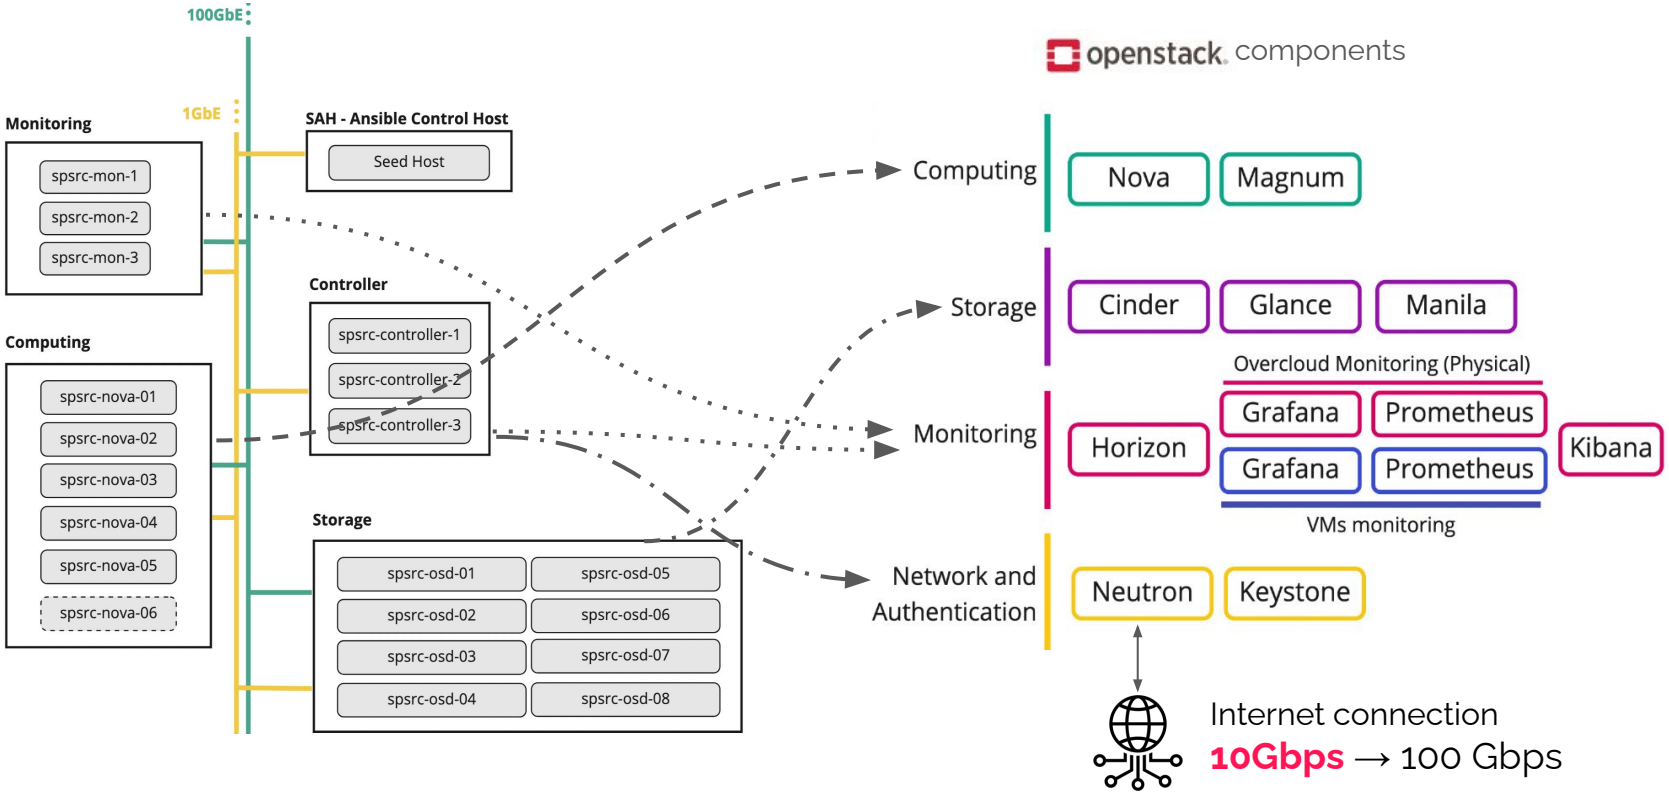

The **esp**SRC computing infrastructure model

OPENSTACK VICTORIA

5 Computing Nodes

- Physical cores: 40, 40, 40, 40, 40 = 200 CPU
- Using overcommit within the OpenStack scheduler:
  - vCPUs: 80, 80, 80, 80, 80 = **400** vCPUs

Wide range of computing flavors:

Name	CPU	RAM	Local HD	Target/Capabilities
spsrc.c2m4	2	4GB	50GB	Light load
spsrc.c4m8	4	8GB	50GB	
spsrc.c8m32	8	32GB	50GB	Multi-purpose
spsrc.c16m64	16	64GB	50GB	
spsrc.c24m194	24	194GB	50GB	
spsrc.c40m1000	40	1000GB	50GB	High load
sdc2.c16m64	16	64GB	100GB	Data challenges
sdc2.c32m128	32	128GB	100GB	

+ SDC,
Dask clusters,
Slurm clusters,
specific purposes,
etc.

The **esp**SRC computing infrastructure model

Octopus (2021) → Quincy (Fall 2023)

8 Storage Nodes managed by Ceph

- Replica 2 → from 1.5 PB to 700 TB
- Block Storage and CephFS

Different pools leveraging disk heterogeneity

- Backups, volumes, metrics, OS's images, shared file systems, RUCIO storage, CADC storage…

The **esp**SRC: Computing services and resources

The **esp**SRC: Computing services and resources

On demand Virtual Machines

- With **Ansible playbooks** to automatise VMs services deployment and configuration.
- All VMs provided with **external public IPs** and FQDN, or entrypoint/header node.
- Only a strict **range of open ports**
- Access protocols: **SSH** and **RDP** (or with Apache Guacamole)
- One user account (PI) with privileges (manage users and install software).
- Storage on demand: from CephFS (Shared Storage) or Block Storage.
- ContainersHub ready-to-use → Singularity
- VMs preloaded with injected radio-astronomy software, services and tools: Individual JupyterHubs, CARTA, etc.
- **Scientific** and IT **support** → ESPSRC User Documentation, a specific Slack channel for daily communication, on-demand meetings and updates.
- Renewal every 6 months.
- Metrics and usage statistics available on Grafana for each project

The **esp**SRC: Computing services and resources

On demand Kubernetes Clusters and Services Hub

- DaskHub:
  Kubernetes cluster for Parallel and Distributed Computing through Python Dask

- ArtifactsHub:
  Kubernetes cluster working as HA containers and artifacts hub

- JupyterHub:
  Kubernetes cluster with several Hubs: ESCAPE-IAM, AMIGA, IAA-Users, RucioHub and Training Schools.

On demand mini Slurm clusters

- MeerKATSlurm: a elastic Slurm cluster to process the MeerKAT pipeline

The **esp**SRC: Computing services and resources

Software distribution and delivery

- **ContainersHub**: a catalogue of Singularity images
    - Containerized Images with software and services for Radio Astronomy: `CASA`, `Kern Suite`, `wsclean`, `ddfacet`,... etc.
    - Mounted as folder (`/mnt/software/containers/`) within each project or cluster.

- **ArtifactsHub**: Containers images (docker, singularity, OCI, etc), Software metadata, Helm charts, etc.

The **esp**SRC: Computing services and resources

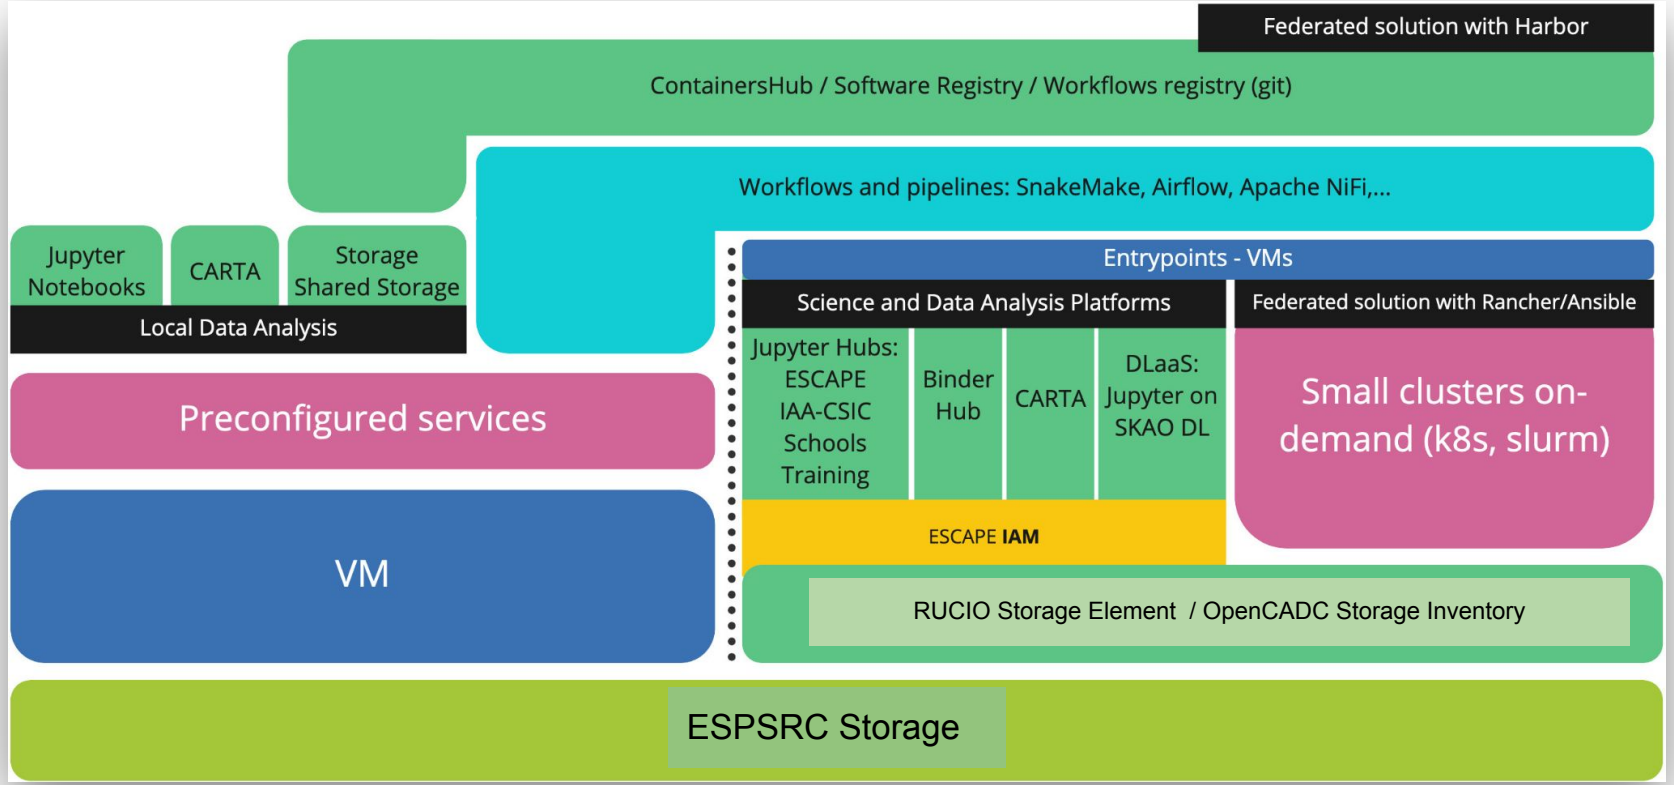

IAA at SRCNet Development

SRCCG, SRC Coordination Group (2016-2018) → LVM invited as external advisor
SRCSC, SRC Steering Committee (2019-) → LVM designated as Spanish representative

SRCSC Working Groups (2020-2022) – Design phase

- WG0 - Architecture
- WG1 - Data Logistics
- WG2 - Operations
- WG3 - Software Processing/Work Flow
- WG4 - Science Archive/VO/FAIR
- WG5 - Compute
- WG6 - Science user engagement

Main outputs
SRC Network requirement gathering (206)
A high-level architecture

IAA at SRCNet Development: Prototypes for the SRCNet

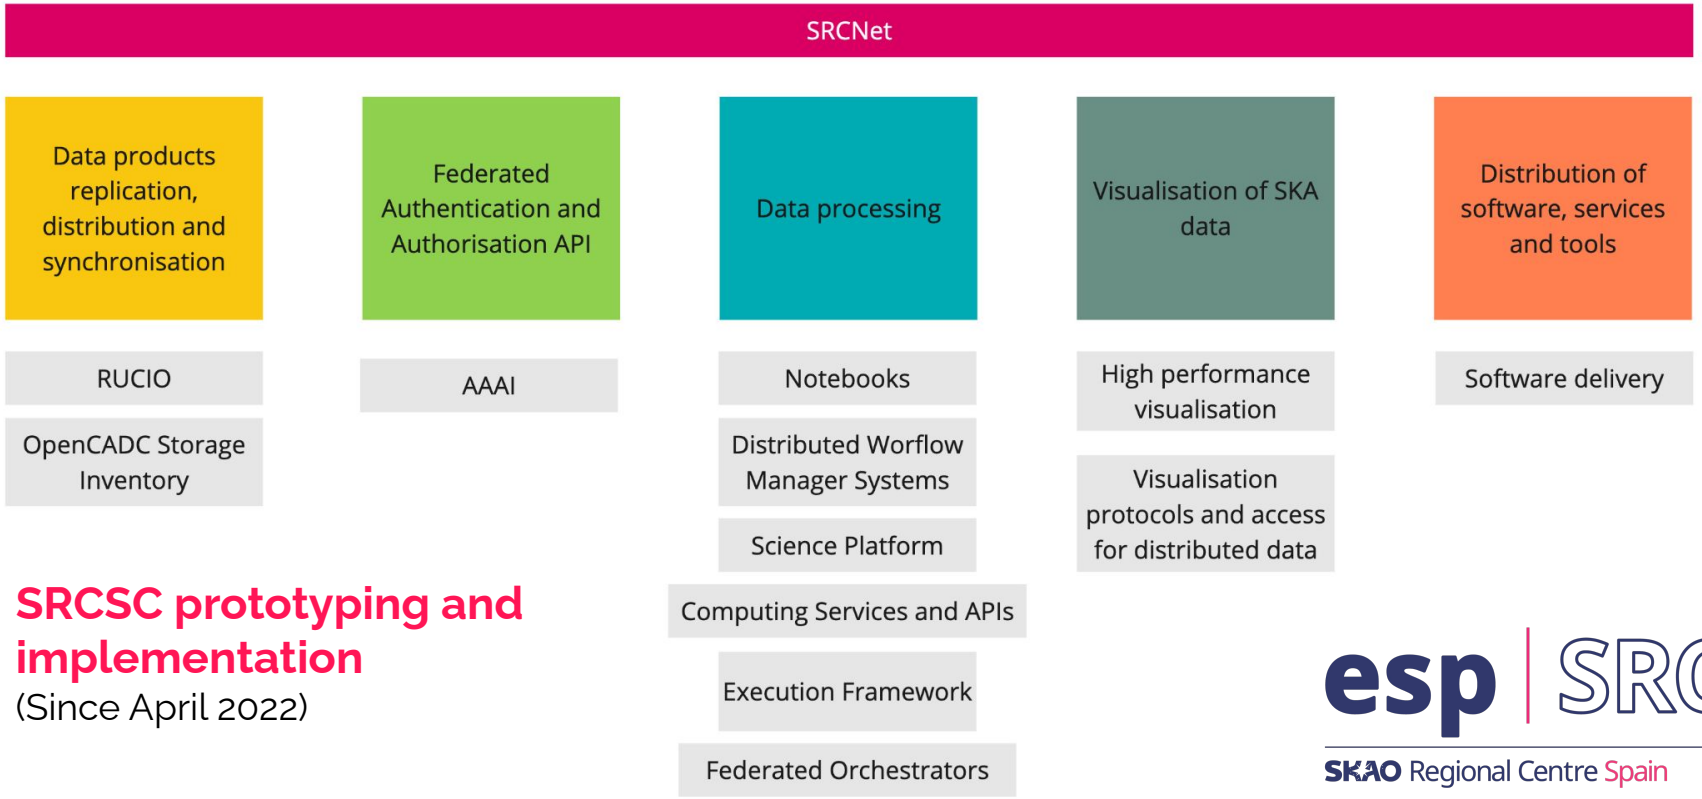

SRCSC prototyping and implementation
(Since April 2022)

esp | SRC
SKAO Regional Centre Spain

How we are building the SRCNet prototypes

SKAO has adopted Agile Methodology and Teams Structure for SRCNet Prototypes

SAFe® PROVIDED BY SCALED AGILE

SAFe Certified

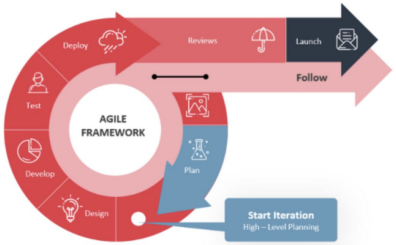

Aim: To provide agility in our development processes, increase efficiency, and deliver value more effectively.

Teams:

- **Stream-aligned** teams are end solution-aligned and are capable of performing all the steps needed to build end-to-end solution value.
- **Complicated subsystem teams** are organized around critical solution subsystems.
- **Platform teams** provide application services and APIs for stream-aligned teams to be able to leverage common platform services.
- **Enabling teams** provides tools, services, and short-term expertise to other teams.

How to build these SRCNet prototypes

9 Agile Teams:

Magenta Team	→ Enabler Team: **Data Management Technologies**
Tangerine Team	→ Enabler Team: **Science Platform Components**
Orange Team	→ Complicated subsystem team: **Visualization**
Olive Team	→ Complicated subsystem team: **HPC services**
Purple Team	→ Enabler Team: **Authentication, Authorisation, Accounting Infrastructure**
Coral Team	→ Stream aligned Team
Blue Lavender	→ Stream aligned Team
Teal A and Teal B	→ Stream aligned Team

Each **Program Increment** lasts 3 months and focuses on **several objectives**. Sprints every 2 weeks.

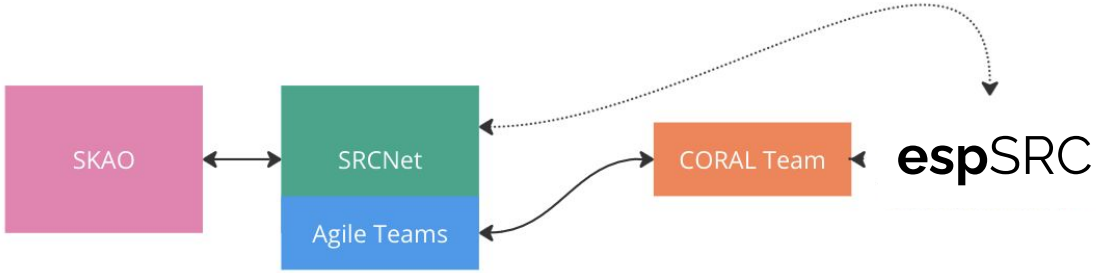

How to build these SRCNet prototypes

9 Agile Teams:

Magenta Team	→ Enabler Team: **Data Management Technologies**
Tangerine Team	→ Enabler Team: **Science Platform Components**
Orange Team	→ Complicated subsystem team: **Visualization**
Olive Team	→ Complicated subsystem team: **HPC services**
Purple Team	→ Enabler Team: **Authentication, Authorisation, Accounting Infrastructure**
Coral Team	→ Stream aligned Team
Blue Lavender	→ Stream aligned Team
Teal A and Teal B	→ Stream aligned Team

Each **Program Increment** lasts 3 months and focuses on **several objectives**. Sprints every 2 weeks.

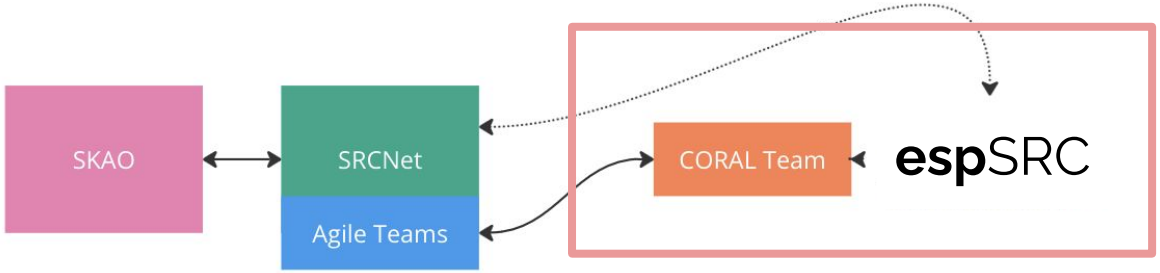

How to build these SRCNet prototypes

CORAL Team @ SRCNet - Stream aligned team with **13** international members

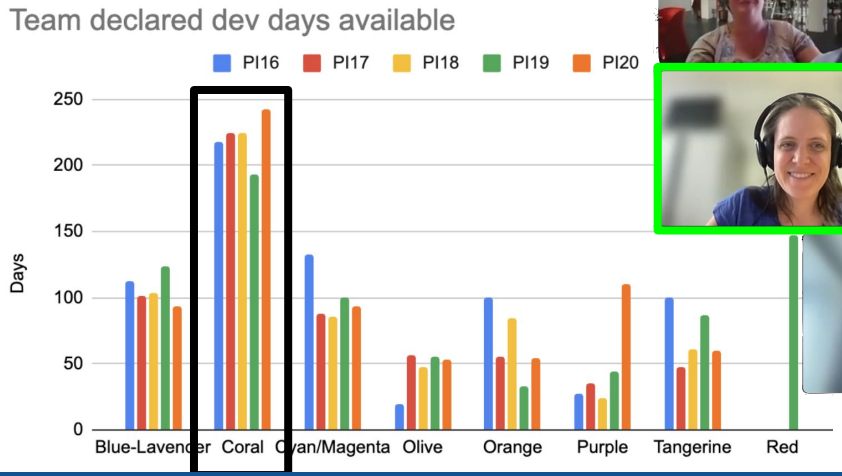

espSRC: Data distribution and Archiving

RUCIO Data Distribution Platform and RUCIO Storage Element (RSE)

A data management system for scientific collaborations

Open-source software - A CERN development

Scalable, robust, and efficient. Supports heterogeneous storage systems within RSEs and globally distributed data centers.

Used by a variety of scientific communities, including high-energy physics, astronomy, and biology

A unified interface across heterogeneous storage and network infrastructures

Monitoring and analytics tools

espSRC: Data distribution and Archiving

Within program Increment #15

Integrate Spanish storage in the prototype 1a testbed - SKAO DataLake

❓ **Who for:** SRCNet Prototype, Other SRC nodes (within Coral or blue-lavender team) that would like to integrate Ceph storage in the Rucio data lake testbed.

🆎 **What**: A data lake instance containing Spanish storage Ceph resources. A recipe for integrating Ceph storage to integrate storage into the Rucio Data Lake. Documentation about the process.

🔎 **Why**: Find out how to integrate in a data lake different storage systems, including Ceph. Increase our knowledge about data lake technologies as Rucio. Identify potential bottleneck/issues adopting data lake technologies as Rucio. Enrich with a new SRC node the data lake testbed.

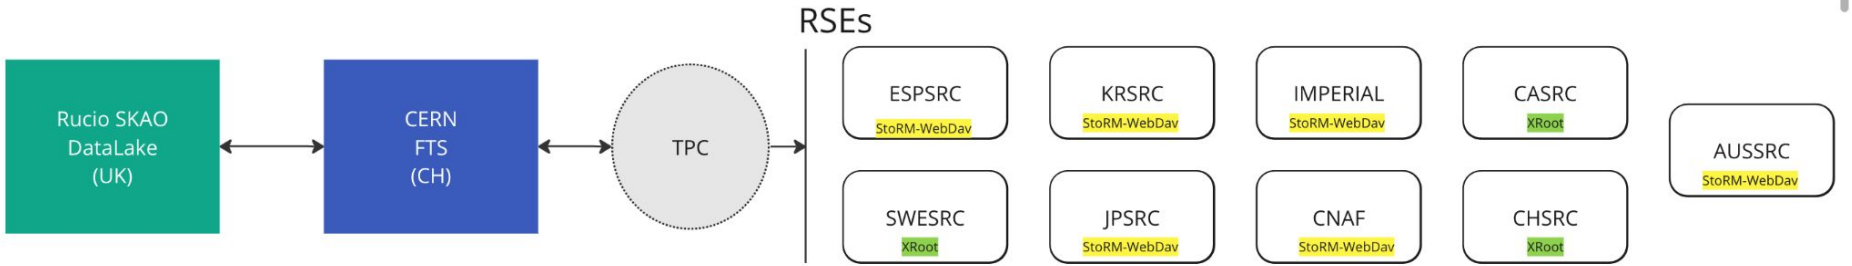

espSRC: Data distribution and Archiving

espSRC Rucio Storage Element (RSE) ready

https://spsrc14.iaa.csic.es:18027/

Storage Resource Manager: StoRM + WebDAV

A&A based on IAM, currently from INAF as AAAI (later with SKAO)

Storage capabilities:

- 10 TB of CephFS storage easily extensible.

RUCIO Network tests by using `rucio-tasks-manager`

- Transfer capabilities, upload & replication, TPC, Tokens, Latency, etc.

espSRC: Data distribution and Archiving

SCIENTIFIC DATA MANAGEMENT

RUCIO

Nodes health

Network tests

Src\Dst	SPSRC_STORM	SESRC_XRD_RBD	KRSRC_STORM	JPSRC_STORM	IMPERIAL	CNSRC_STORM	CNAF	CHSRC_XRD	CASRC_XRD	AUSSRC_STORM
STFC_STORM	0%	0%	0%	0%	0%	100%	0%	0%	0%	100%
SPSRC_STORM	-	0%	0%	0%	0%	100%	0%	0%	0%	100%
SESRC_XRD_RBD	0%	-	1%	0%	0%	100%	0%	0%	0%	1%
KRSRC_STORM	1%	0%	-	0%	0%	100%	0%	0%	0%	100%
JPSRC_STORM	0%	0%	0%	-	0%	100%	0%	0%	0%	100%
IMPERIAL	0%	0%	0%	0%	-	100%	0%	0%	0%	100%
CNAF	0%	0%	0%	0%	0%	100%	-	0%	0%	100%
CHSRC_XRD	0%	0%	1%	0%	0%	100%	0%	-	0%	100%
CASRC_XRD	0%	0%	0%	0%	0%	100%	0%	0%	-	100%
AUSSRC_STORM	0%	0%	0%	0%	0%	100%	0%	0%	1%	-

espSRC: Data distribution and Archiving

SCIENTIFIC DATA MANAGEMENT

RUCIO

Nodes health

Network tests

file size \ RSE	SPSRC_STORM	KRSRC_STORM	JPSRC_STORM	IMPERIAL	CNAF	CHSRC_XRD	CASRC_XRD
200 MB	28.4 MB/s	6.58 MB/s	9.96 MB/s	30.0 MB/s	32.8 MB/s	39.4 MB/s	12.1 MB/s
500 MB	42.1 MB/s	9.05 MB/s	14.1 MB/s	42.7 MB/s	49.0 MB/s	52.1 MB/s	18.0 MB/s
1 GB	46.4 MB/s	9.42 MB/s	14.2 MB/s	48.0 MB/s	48.4 MB/s	50.4 MB/s	17.0 MB/s
1.50 GB	52.0 MB/s	9.53 MB/s	14.4 MB/s	43.2 MB/s	65.2 MB/s	48.7 MB/s	16.8 MB/s
2 GB	53.1 MB/s	9.60 MB/s	14.3 MB/s	39.3 MB/s	54.6 MB/s	48.9 MB/s	16.7 MB/s
3 GB	44.7 MB/s	9.62 MB/s	14.2 MB/s	41.7 MB/s	56.8 MB/s	49.6 MB/s	16.5 MB/s
4 GB	40.8 MB/s	9.67 MB/s	14.4 MB/s	41.0 MB/s	54.4 MB/s	49.8 MB/s	16.4 MB/s
Mean	43.9 MB/s	9.07 MB/s	13.6 MB/s	40.9 MB/s	51.6 MB/s	48.4 MB/s	16.2 MB/s

espSRC: Data distribution and Archiving

Distributed Archive Data Management System developed by the Canadian Astronomy Data Centre (CADC) → OpenSource

Designed to manage archival file storage for a science data archive:

Data and Metadata

Manage two entities:

- **Local sites** → back-end storage (for files), a database with the local inventory, a **files** service, and a **metadata** service
- **Global inventory** → database with inventory of all files from all storage sites, a **locator** service, a **metadata** service, a **metadata-sync** process for each storage site, and a periodic **metadata-validate** process for each storage site

A&A based on certificates and IAM tokens.
Highly flexible and scalable. Based on containers and microservices.
Local sites and Global inventory enabled to run on Kubernetes.

espSRC: Data distribution and Archiving

OpenCADC

SI

Storage Inventory

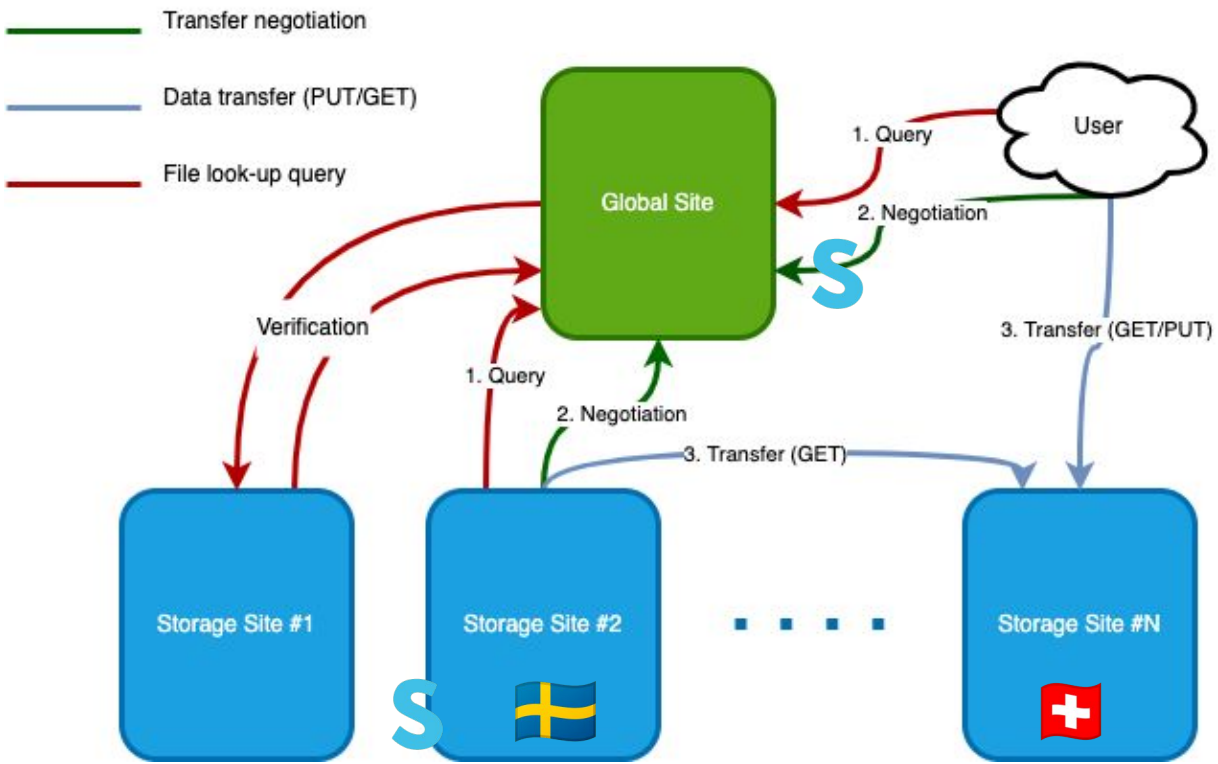

espSRC: Data distribution and Archiving

Local site services:

- https://spsrc29.iaa.csic.es/minoc/ → Files
- https://spsrc29.iaa.csic.es/luskan/ → Metadata
- https://spsrc29.iaa.csic.es/argus → Table Access Protocol (TAP) for CAOM
- https://spsrc29.iaa.csic.es/bifrost/ → Datalink Service for CAOM

Global inventory:

- https://spsrc27.iaa.csic.es/luskan/ → Metadata
- https://spsrc27.iaa.csic.es/raven/ → FIle location
- https://spsrc27.iaa.csic.es/baldur/ → Permissions
- https://spsrc27.iaa.csic.es/reg/ → Registry

~14.6 TB

Global registry

- https://spsrc31.iaa.csic.es/torkeep → Metadata Repository
- https://spsrc31.iaa.csic.es/argus → Table Access Protocol (TAP) for CAOM

The Mini-SRC

Objective
Deploy an end-to-end multi-SRC node Demonstrator with data management, science archive, science platform, user collaboration spaces and execution planning with integrated A&A and interoperability across all sites.

Mayor components

- Integrated A&A using prototype **IAM** service
- **Data management** and **dissemination,** and Science archive (metadata management and data access, UI, VO services)
- **Science platform** (user container images, notebooks, visualization, desktop sessions, UI, VO services)
- User **collaboration spaces** (mountable storage, UI, VO services)
- **Execution planning** (code to data, data to code, data and code to available resource)
- **Credential delegation** service.

Mini-SRCNet nodes

Science, Development and Training projects

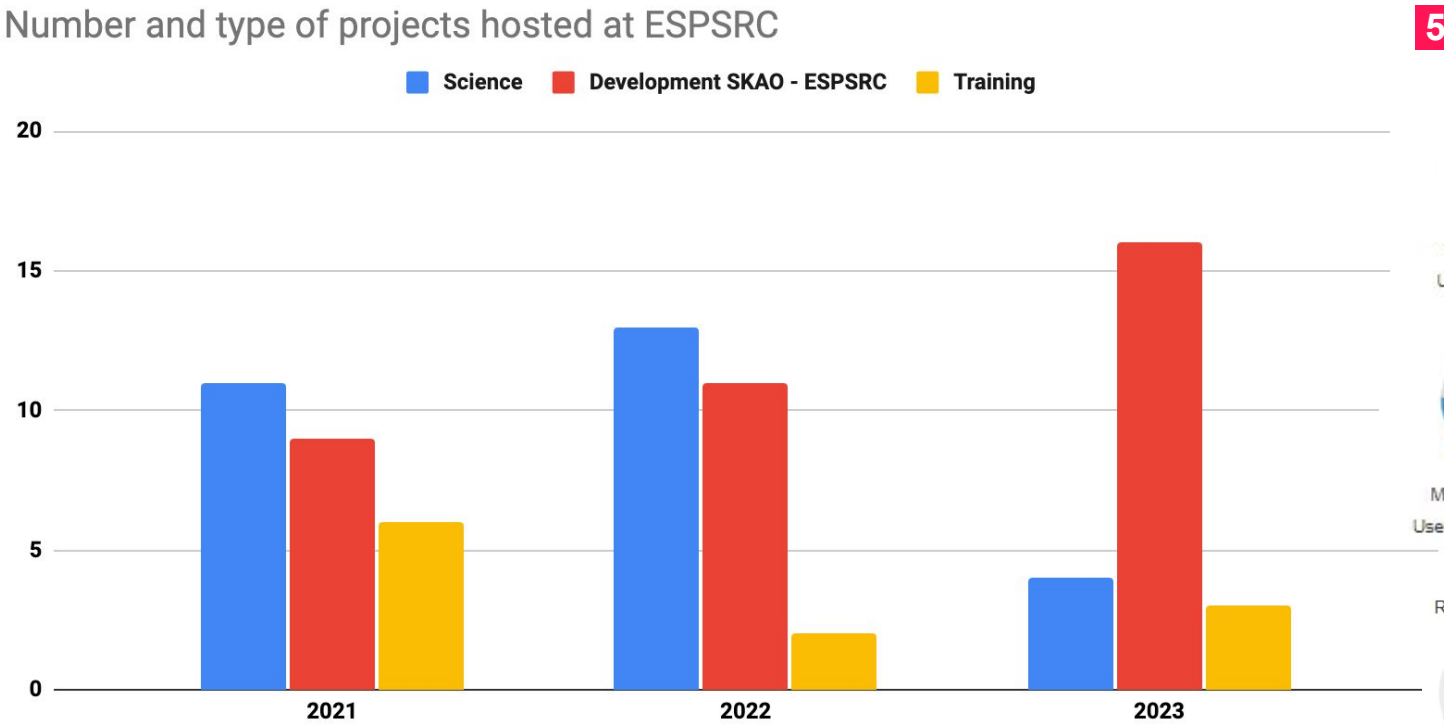

52 VMs

VCPU Usage

Used 436 of 240

Memory Usage

Used 1.9TB of 2.5TB

Raw Capacity

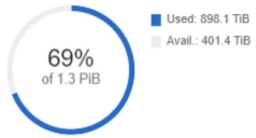

60 projects: 23 Research, 25 Development, 12 Training (incl. SKA Data Challenges).

Science, Development and Training projects

Projects

Projects by Instrument

SKA Data Challenges

Current load

Research	Development	Training
Study of Star-planet interaction MeerKat	CARACAL - Data reduction pipeline	LOFAR School 2021
Galactic Radial Profiles	SoFIAX Services	SOMACHINE 2020 *
Lirgi Luminous Infra-Red Galaxy Inventory (LIRGI)	AusSRC - DaCHS VO - MultiMaster	SOMACHINE 2021 *
LOFAR - Study of Exoplanets	NextCloud Docs & Repository	JupyterLab for IAA-CSIC *
Perseus - Perseus Radio Observations	Network Test	SKA Training on Containers
Hydra - Gas and star formation in the Hydra Cluster	SlicerAstro deployment	JupyterLab for PySnacks Workshop *
APERTIF - Apertif: HI Source Finding	Singularity builder and containers	SKA Data Challenges: Ghydrann
AMIGA-J - Angular momentum of isolated AMIGA galaxies	Openstack services verification	SKA Data Challenges: HI-friends
Proxima-ALMA - Searching for dust components in Proxima Centauri planetary system with ALMA	JupyterLab for WALLABY *	Radio Master - Course on Radio Astronomy
ALMA.W75N - Analysis of massive star-forming region W75N	BinderHub *	
Grain growth in protoplanetary disks in Oph A cluster	CARTAHub *	
Numerical simulations of the exospace weather around the planet Proxima b using the 3D MHD	Monitoring infrastructure	
3D modelling of the solar corona	DLaaS	
Differential photometry with LEMON	Harbor	
CHANG-ES data reduction	Tangerine Resources	
	RUCIO RSE	

Distributed as:
- Standalone VMs
- Clusters of VMs
- K8s clusters

Science, Development and Training projects

Projects

Projects by Instrument

SKA Data Challenges

Current load

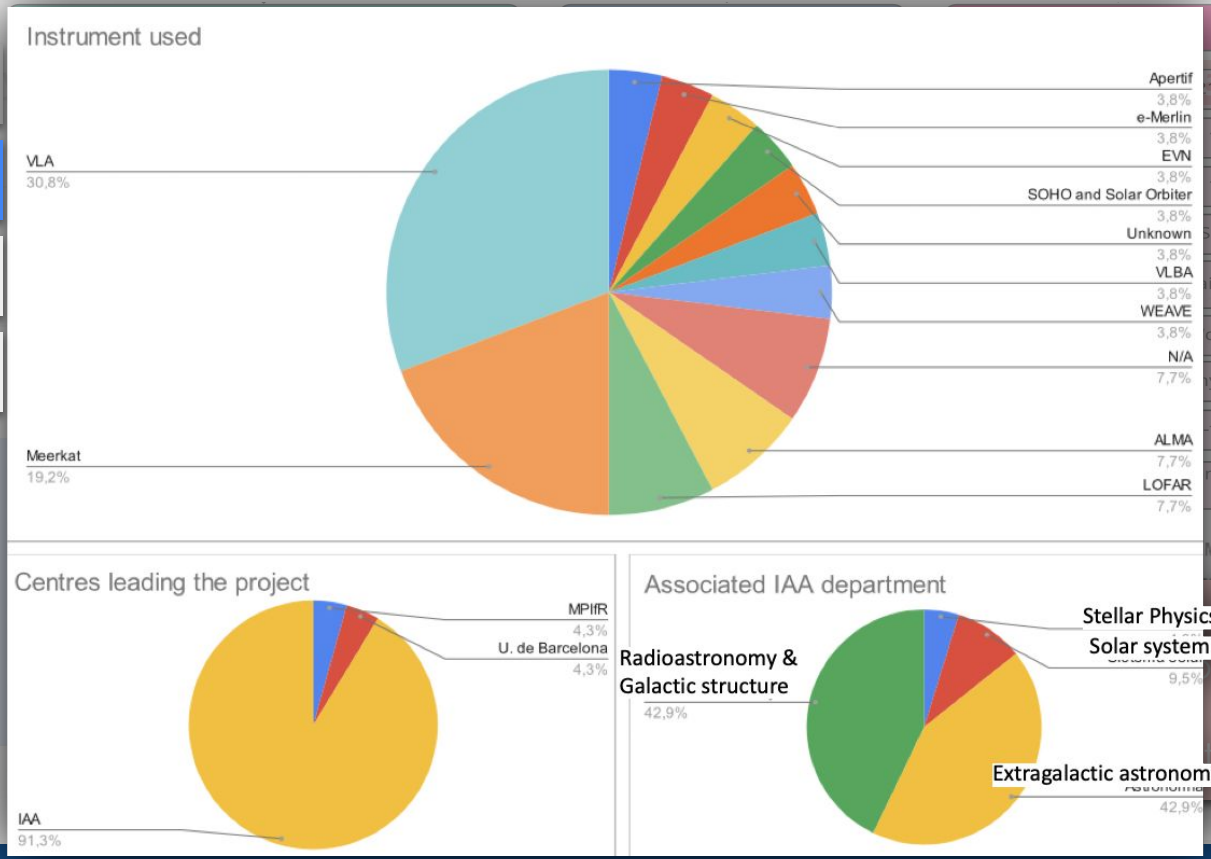

Science, Development and Training projects

Projects

Projects by Instrument

SKA Data Challenges

Current load

CESGA
GALICIA SUPERCOMPUTING CENTER

Supporting SDC3 teams (with Cloud Computing resources)

IFCA
Instituto de Física de Cantabria

Team participating in the **SKA Data Challenge 3, supported by espSRC**

Science, Development and Training projects

Projects

Projects by Instrument

SKA Data Challenges

Current load

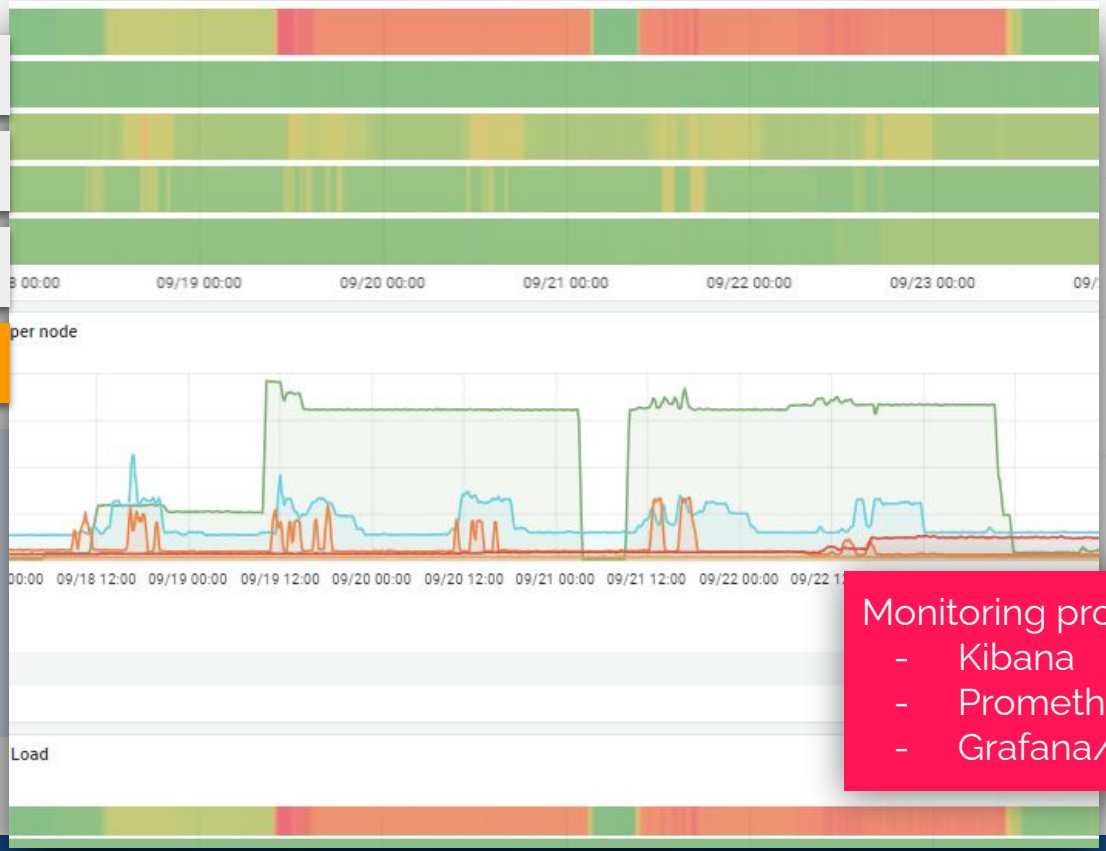

52 VMs_

VCPU Usage
Used 436 of 240

Memory Usage

Monitoring projects and ESPSRC:

- Kibana
- Prometheus
- Grafana/Zabbix

Open Science & Collaborative environment

https://github.com/espsrc/

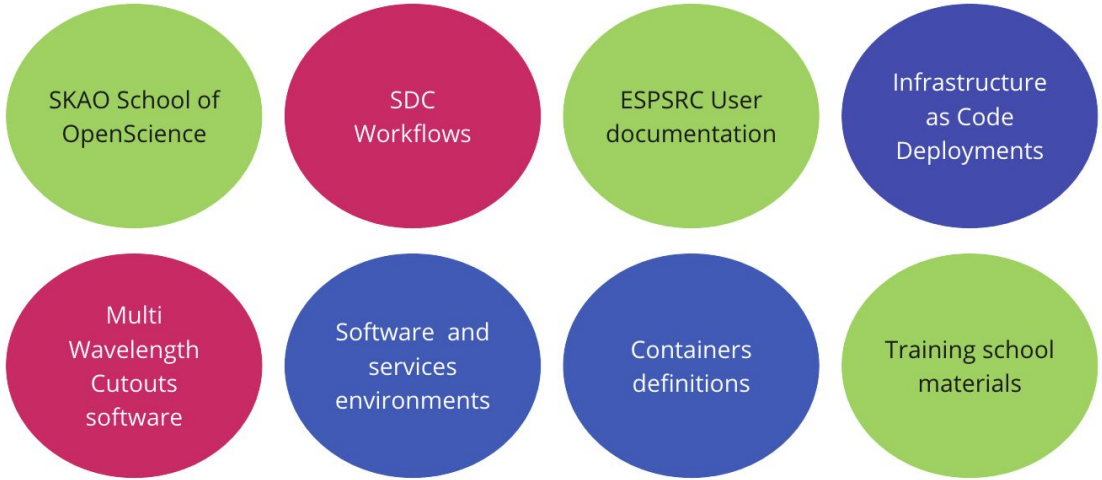

Open Science & Collaborative environment

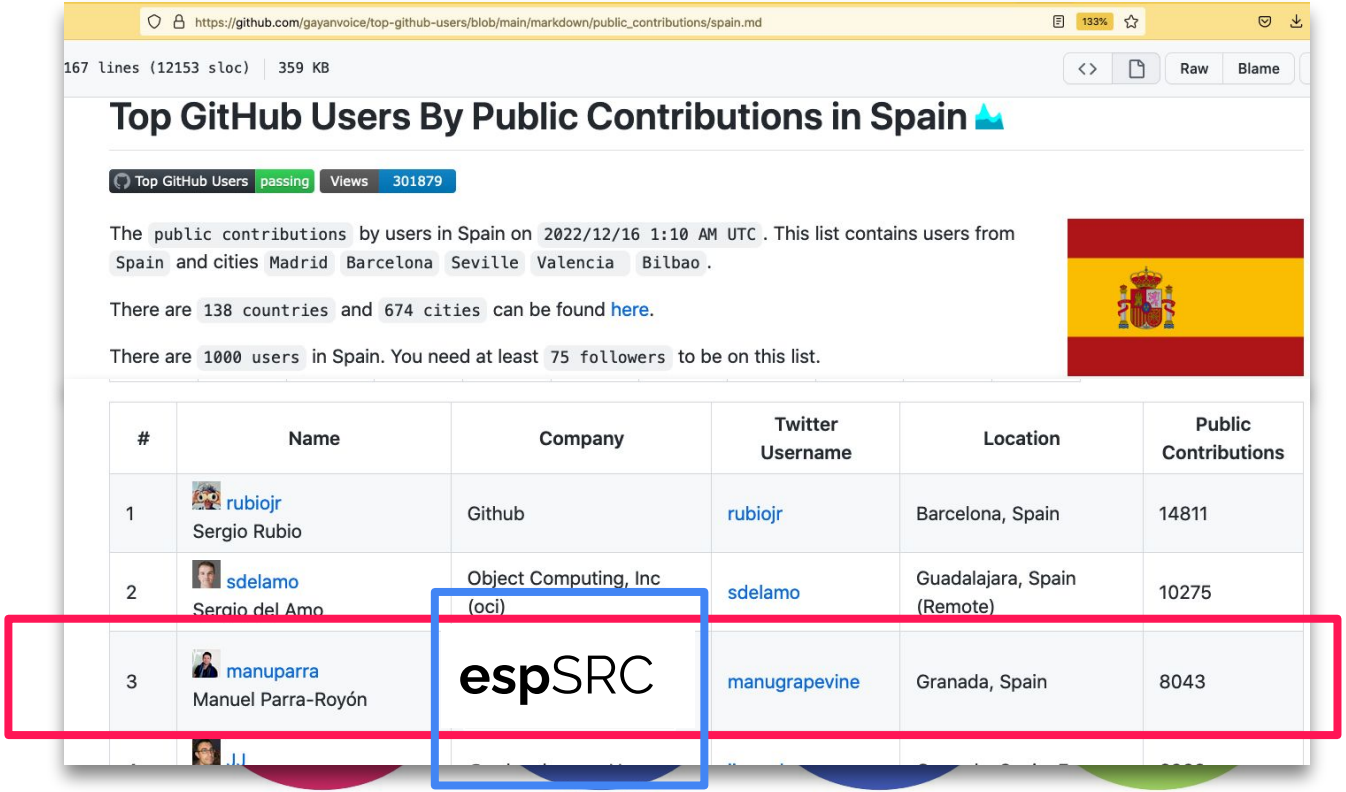

Conclusions

The SKAO project represents an unparalleled computational challenge in the exploration of the cosmos.

SRCNet will be the backbone that will enable the processing and distribution of data to the global scientific community. The building blocks of SRCNet are supported by the development of cutting-edge prototypes developed by international teams.

National infrastructures are essential for the success of the prototypes and the team working on the prototypes: The ESPSRC offers a powerful and flexible platform that encompasses a wide range of resources and services

ESPSRC fosters a transparent and efficient collaboration in alignment with the principles of FAIR and Open Science.

ESPSRC is providing first-class support service to the SKAO community and beyond

References

[1] SKA Regional Centre Prototype at IAA-CSIC: building an Open Science platform based on cloud services. Sánchez-Expósito, S.; Luna, Sebastián; Garrido, Julián, et al. ADASS 2020. Nov 2020.

[2] Toward a Spanish SKA Regional Centre fully engaged with open science. Julián Garrido, Laura Darriba, Susana Sánchez-Expósito et al. Journal of Astronomical Telescopes, Instruments, and Systems, Vol. 8, Issue 1, 011004 (November 2021). https://doi.org/10.1117/1.JATIS.8.1.011004

[2] Asymmetric distribution of data products from WALLABY, an SKA precursor neutral hydrogen survey. Manuel Parra-Royón, Austin Shen, Tristan Reynolds, et al. ADASSXXI. Nov 2021. Preprint: https://arxiv.org/abs/2303.11670

[3] Cloud services at the Spanish SKA Regional Centre prototype (SPSRC) and what could be used by the SRCNet. Susana Sánchez Expósito, Lourdes Verdes-Montenegro, Julian Garrido, et al. SRC WG5 Workshop: Native Cloud Computing and Software Defined Clusters. Mar 2023.

[4] A fully-reproducible workflow for the SKA Data Challenge 2 HI-FRIENDS solution. Javier Moldón; Laura Darriba; Lourdes Verdes-Montenegro, et al. Reproducibility and Open Science in Astronomy, May 2022.

[5] A Spanish Prototype of a SKA Regional Centre. Javier Moldón. SEA Conference 2022, Sep 2022.

[6] Integration of storage endpoints into a Rucio data lake, as an activity to prototype a SKA Regional Centres Network. M Parra-Royón, J Sánchez-Castañeda, J Garrido, S Sánchez-Expósito et al.. ADASS XXII. Nov, 2022. Preprint: https://arxiv.org/abs/2303.07524

[7] Open Science and green computing to support sustainability at the SKA Observatory and its Regional Centres. Julián Garrido, Lourdes Verdes-Montenegro, Susana Sánchez. Astronomers for planet earth Symposium 2022. Nov, 2022.

[8] An approach to provide serverless scientific pipelines within the context of SKA. Manuel Parra-Royón (presenter), Carlos Ríos-Monje, Javier Moldón, et al. IWSG 2023, Tubingen, Germany (Jun 2023).

[9] Empowering SKA Data Challenges: A homogeneous platform for enhanced collaboration and scalability fully aligned with Open Science. Manuel Parra-Royón, James Collinson, Jesús Sánchez-Castañeda, Pablo Llopis-Sanmillan, Carolina Lindqvist, Susana Sánchez-Expósito, Laura Darriba and Lourdes Verdes-Montenegro. ADASSXXIII. Nov 2023.

The ESPSRC's approach to an ecosystem of services for science support within the SKA Regional Centres Network

Manuel Parra-Royón, Lourdes Verdes-Montenegro, Susana Sánchez-Expósito, Julian Garrido, Laura Darriba, Javier Moldón, María Ángeles Mendoza, Jesús Sánchez-Castañeda

Instituto de Astrofísica de Andalucía - CSIC

IBERGRID 2023
Benasque, Huesca, Spain

With the financial support from a) and b):

a) 

b) The grant **CEX2021-001131-S** funded by MCIN/AEI/10.13039/501100011033 and from the grant **DOC01497** funded by the Economic Transformation, Industry, Knowledge and Universities Council of the Regional Government of Andalusia and by the Operational Program ESF Andalucía 2014-2020